An Evaluation of the Security of MULTI-S01

David Wagner

December 17, 2001

Executive Summary

This report presents the report of a very limited examination the security of the MULTI-S01 cipher. No attacks on the cipher were found.

1 Introduction

MULTI-S01 is a cryptographic algorithm proposed in the CRYPTREC standards effort that uses the PANAMA stream cipher to provide both encryption and message authentication services. In this report, I examine the security of MULTI-S01.

MULTI-S01 uses PANAMA in the following way. It takes as input a 256-bit secret key K and a 256-bit IV ("deviation parameter") Q. The PANAMA internal memory is zeroed. Then, we execute two PUSH operations with K and Q to key the cipher. Finally, we execute a series of PULL operations to produce a long stream of pseudorandom output $\text{PANAMA}(K,Q)$. The pseudorandom stream is used to encrypt and authenticate the message, and we transmit the result along with the IV Q in the clear. A separate IV must be used for each message.

2 Chosen-IV attacks on MULTI-S01

Because of the way that PANAMA is used, it appears that one of the most promising avenues for analysis of MULTI-S01 is the chosen-IV attack. In a chosen-ciphertext threat model, the attacker has an opportunity to select both the IV Q and the ciphertext and obtain the result of decrypting this ciphertext (if it is authentic).

Thus, it seems to make sense to consider the security of PANAMA in the following sense. The adversary can choose many Q values and obtain the pseudorandom stream $\text{PANAMA}(K,Q)$ produced by PANAMA under key K and IV Q. The adversary wins if she can distinguish the resulting pseudorandom streams from a truly random source. If PANAMA is secure in this setting, then MULTI-S01 will be secure, too.

In other words, the proposed evaluation criteria is as follows:

> *Security requirement.* The function $F_K(Q) = \text{PANAMA}(K,Q)$ should form a secure pseudorandom function.

If PANAMA does form a secure pseudorandom function in this sense, MULTI-S01 will be secure. If PANAMA does not form a secure pseudorandom function, it is possible that there may be some attack on MULTI-S01, and in this case it would appear to be prudent to abandon use of MULTI-S01. As a result, we obtain a simple condition that expresses the properties required of PANAMA.

What is currently known about the security of PANAMA as a pseudorandom function? There have been no published attacks on this mode of usage of PANAMA, and I have not been able to find any weaknesses, either. Nonetheless, I will survey some of the ways that cryptanalysts might try to find attacks in PANAMA.

2.1 Chosen-IV collision attacks

PANAMA has a very large internal internal memory, consisting of two parts: the *state* (544 bits), and the *buffer* (8192 bits). However, if the cryptanalyst can somehow find two IV's Q, Q' so that the internal memory after pushing K and Q is the same as the internal memory after pushing K and Q', then the cryptanalyst will have succeeded in

breaking PANAMA, because in this case the two output streams will be the same. I will call this a *chosen-IV collision attack*.

Consequently, it seems important to study the security of PANAMA against chosen-IV collision attacks. Fortunately, we can say the following:

Theorem 1. *PANAMA (as used in MULTI-S01) is secure against chosen-IV collision attacks.*

Proof. Let $M_K(Q)$ represent the value of the intermediate memory of PANAMA after pushing K and Q. Then the function $M_K : \{0,1\}^{256} \to \{0,1\}^{8736}$ is one-to-one for all $K \in \{0,1\}^{256}$, as may be readily verified by examining the structure of PANAMA: the IV Q is XORed against the memory, and then only bijective operations are performed as we execute pulls. In other words, for all K, Q, Q' with $Q \neq Q'$, we have $M_K(Q) \neq M_K(Q')$, and thus there is no possibility of causing a collision in the internal memory of PANAMA using chosen-IV attacks. □

The proof of the above theorem relies crucially on the fact that the IV in MULTI-S01 is only 256 bits long, and hence is fed into PANAMA all in a single push operation.

If the IV were 512 bits long, and hence spread across two consecutive push operations, I believe it might be possible to choose a pair of IV's that yield collisions with high probability by using differential techniques. The reason is that changing a single bit in the first 256 bits of the IV will spread to only three bit positions in the intermediate state after a single push operation, with probability $1/4$, and if we carefully choose that bit position, we should be able to ensure that the next 256 bits of IV are chosen in such a way to cancel out this difference. Differences in the buffer will soon lead to avalanche, but it may be possible to control this enough to distinguish PANAMA from random. Since this attack is clearly not possible on MULTI-S01, I did not study it in detail any further.

Rijmen, et al., have shown that it is possible to find collisions for PANAMA with complexity 2^{82} [2]. Their analysis uses a differential analysis of the PANAMA transformations to obtain a collision in the internal memory after pushing a number of message words. As a consequence, one could imagine that it might be possible to exhibit a similar attack on MULTI-S01, replacing their chosen messages with chosen IV's. However, this approach does not seem likely to lead to a collision attack on MULTI-S01, because they rely heavily on the fact that they can choose long messages that will be spread across many push operations.

2.2 Chosen-IV differential attacks

More generally, we might consider differential attacks on PANAMA where the intermediate differences in the internal memory of PANAMA are predicted after we inject some chosen difference into the IV Q. In other words, we can ask whether we can find some differential $\Delta Q \to \Delta M$ for the push operation (where ΔM represents the difference in the internal memory) along with some differential $\Delta M \to \Delta M'$ for the 33 blank pull operations. If we could find such differentials of sufficiently high probability, we might be able to distinguish PANAMA from random in a chosen-IV attack.

Note that this is a generalization of the chosen-IV collision attack. In a collision attack, we seek to make $\Delta S = 0$, but in a chosen-IV differential attack, we only ask that ΔS be predictable.

I do not know whether there exist any good differentials for PANAMA that might be usable in a chosen-IV differential attack. I tried to find good differentials for this purpose, and I failed. I was unable to find any differentials of non-negligible probability, because differences avalanche very rapidly. It does not seem easy to find good differentials for 33 blank pulls: they introduce a great deal of non-linearity and avalanche.

Rijmen's analysis shows that there exist some fairly high-probability differentials for PANAMA, but because their analysis relied on ability to specify chosen differences at the input to a number of push operations, it is not clear whether this will extend to analysis of MULTI-S01. Nonetheless, their style of analysis seems to set the direction for future evaluation of PANAMA and raises at least the potential that there may be some way to control the spread of differences in PANAMA long enough to see them appear in the output.

2.3 Chosen-IV related-key differential attacks

In a generalization of this threat model, we might allow the adversary the ability to choose not only differences ΔQ to appear in the IV, but also differences ΔK for the key. In other words, the attacker can specify ΔQ and ΔK freely and obtain PANAMA$(K \oplus \Delta K, Q \oplus \Delta Q)$, but the key K remains secret throughout.

This threat model is closely analogous to expanding the IV to 512 bits, and my earlier comments regarding that variant of MULTI-S01 apply here as well. The conclusion is that there seems to be a very real chance that there could be an effective chosen-IV related-key differential attack on MULTI-S01.

However, I do not think this style of attack is very relevant. The related-key threat model appears to be primarily of theoretical interest, and in most cases there is no way for attackers to make related-key queries. As a result, I did not study this threat model in depth. If security against related-key attacks is considered important, then this aspect should be investigated further.

2.4 Analysis of simpler variants of MULTI-S01

Suppose that we omit the 33 blank pulls used before generating pseudorandom output with PANAMA. Then it is very easy to obtain a distinguishing attack on this variant. In particular, the last push does not affect the first 256 bits of keystream output, and so the first 256 bits of keystream will be independent of the IV.

Suppose we next consider a variant with exactly 1 blank pull. This variant is also easy to break. For example , we have the differential

$$(0,0,0,0,0,1,0,0,0,0,\ldots,0) \xrightarrow{\rho} (0,0,0,0,16,0,0,16,16,0,\ldots,0) \qquad \text{with prob. } 1/4$$

for the state-update function. (There are also many others.) Consequently, if we insert the difference $\Delta Q = (0,0,0,0,1,0,0,0)$ into the IV, the next pull operation will yield

256 bits of keystream output that have the difference $(0,\ldots,0)$ with probability $1/4$. This is an effective distinguishing attack.

This analysis may be extended further. In fact, with up to about 14 blank pulls, we obtain a key complementation property. Suppose we consider applying the difference $\Delta K = (0,0,0,0,1,0,0,0)$ to the key, as a thought experiment, and simultaneously apply the difference $\Delta Q = (0,0,0,16,0,0,16,16)$ to the IV. With probability 1/4, we will have a collision in the state (but not the buffer) after the two pushes using K and Q. Then after three blank pulls these differences are re-introduced into the state from the buffer, but they get immediately cancelled out with probability $1/4$ again. After up to 14 blank pulls, we still have a collision in the state (but not the buffer), and this suffices to ensure that the first 256 bits of keystream will be equal in both encryptions with probability $1/16$.

This is related-key differential of probability $1/16$, and given our earlier comments about the inapplicability of related-key attacks, it may sound uninteresting. However, it can also be viewed as a key-complementation property that holds with probability $1/16$: if we obtain multiple pairs of encryptions with related IV's, then the value of a single key bit does not affect the first 256 bits of output for $1/16$ of these pairs of encryptions. In short, with probability $1/16$ one bit of key material is ignored. Similar observations apply to many other bits of the key.

This probabilistic key-complementation property might permit an attacker to speed up exhaustive keysearch somewhat. This is unlikely to be a threat if we use uniformly distributed 256-bit keys, but it might be a serious threat if the key is chosen from a space with small entropy (e.g., as a passphrase, or using a faulty random number generator).

It is conceivable that it might be possible to set up a similar probabilistic key-complementation property that survives all 33 blank pulls. The tricky part is that when the difference is introduced into the state the third time (after about 15 blank pulls), it is introduced in a different form (namely, into the opposite half of the state), and both forms must have high probability if we are to obtain a useful key-complementation property. I was not able to find a good probabilistic key-complementation property for the full MULTI-S01, but I believe there is a possibility of finding such a property with further study.

3 General comments

3.1 On the specification of MULTI-S01

A comment on the MULTI-S01 specification also seems in order. The MULTI-S01 specification actually does not specify the use of 33 blank pulls: it says to use two push operations to inject K and Q, and then to obtain the keystream using pull operations, but does not mention anything about blank pulls. It seems likely that the designers intended to mandate the use of 33 blank pulls between these two steps, just as is done in ordinary PANAMA, but this does not appear to have been stated anywhere in the specification of MULTI-S01. As I have made clear earlier, MULTI-S01 is insecure if these blank pulls are omitted. This introduces the risk that a implementor who is unaware of this pitfall might follow the specification literally and omit the blank pulls,

thereby obtaining a cryptographic algorithm with a very serious weakness. I note that the reference implementation of MULTI-S01 avoids this pitfall, but I am concerned that independent implementations might be at risk. Therefore, I suggest the specification should be augmented to make the need for the blank pulls absolutely clear.

3.2 On equivalent keys

An earlier security evaluation raised questions about the risk of equivalent keys [1]. Let (A,B,S) denote the pseudorandom values output by PANAMA under key K and IV Q. The concern expressed is that, for any plaintext/ciphertext pair, there are 2^{64} values for (A,B,S) consistent with this pair, and one of the endpoints might use this property to repudiate her presence in MULTI-S01-encrypted session.

This issue does not seem to me to be a problem. First, because the equivalent values of (A,B,S) can be easily identified, any third party can readily detect this sort of falsification. Second, the values (A,B,S) are not the key material: the real key is K, and there is no reason to expect that there will be equivalent keys K,K' both consistent with a single plaintext/ciphertext pair. Third, and most importantly, repudiation is always possible in symmetric-key cryptosystems, and no other symmetric-key encryption algorithm defends against this sort of attack, either.

Therefore, I believe that this concern is not an issue, and does not require any further study.

3.3 On provable security and a large period

An earlier security evaluation pointed out that there is no proof that the period or cycle structure of PANAMA is favorable [1]. Of course, if PANAMA has a small period, then MULTI-S01 will be insecure.

This is correct, but it does not seem to me to be a problem. The same can be said about every other mode of operation, including (for example) widely-accepted systems such as Triple-DES-OFB mode and AES-CTR mode. No proofs are known, but it is widely viewed to be extremely probable that these modes have a very large period.

Moreover, the security requirement articulated earlier already captures the possibility of small-period attacks and much more. If PANAMA forms a secure pseudorandom function, then it has a very long period and excellent cycle structure. In other words, there seems to be little reason to worry specifically about the period of PANAMA; in my opinion, it makes more sense to focus attention on evaluation of whether PANAMA meets the conditions needed to be a secure pseudorandom function.

4 Conclusions

After a limited-time analysis, I was not able to find any weaknesses in PANAMA or MULTI-S01. No serious attacks on PANAMA are known in the open community, and as far as our current knowledge is concerned, PANAMA appears to be adequate for the required purpose. The excellent performance of PANAMA also makes it very attractive. However, PANAMA does not appear to provide quite the same level of assurance

of security as offered by more widely-accepted constructions, such as Triple-DES or AES. I would feel more comfortable if PANAMA had received more scrutiny from the cryptographic community, and it would appear to be prudent to devote further research into the security of PANAMA before placing a great deal of trust in its security. As a consequence, for high-security applications, it may be preferable to stick to accepted standards such as Triple-DES and AES, where feasible.

4.1 Caveats

The results of this evaluation should be interpreted with care. This is a limited time evaluation by a single cryptographer. It is widely accepted in the cryptographic community that acquiring enough confidence in a new cipher usually requires years of scrutiny by the cryptographic community at large. While this report does give some very limited evidence for the security of MULTI-S01, it is no substitute for years of public study.

References

[1] Anonymous, "Evaluation of MULTI-S01," confidential evaluation, 38 pages, January 17, 2001.

[2] Vincent Rijmen, Bart Van Rompay, Bart Preneel, and Joos Vandewalle, "Producing Collisions for PANAMA," *Fast Software Encryption 2001*, Springer-Verlag.